

SPECIFICATION

Electronic Version 1.2.8

Stylesheet Version 1.0

Method and System for Optimizing Leaf Comparisons from a Tree Search

Field of the Invention

[0001] The present invention relates to control structures for tree searches in embedded processing systems.

Background of the Invention

[0002] Processing system designers continually seek new ways to improve device performance. While processing speeds continue to increase, the latency imposed by memory access times imposes operating delays. In systems-on-a-chip/embedded systems, efforts to avoid such latency issues have included utilizing local memory in the form of SRAM (static random access memory) on-chip. However, cost and size limitations reduce the effectiveness of the use of SRAM on-chip for some processing environments.

[0003] For example, currently in network environments, network switches are being used to perform more complex operations than simple packet forwarding. Network processors are being developed to provide for more complex processing in network routers, while maintaining flexibility to accommodate changes and enhancements to the functionality provided by the routers, as techniques and protocols evolve. As with most any form of processors, these network processors also face challenges in terms of memory utilization, particularly due to the need to handle a vast array of network traffic.

[0004] In embedded processing systems, such as network processors, off-chip/external DRAM (dynamic random access memory) is an option that is often chosen due to its

lower cost, as compared with SRAM. Thus, while potentially most cost effective, the use of external DRAM introduces a performance penalty in the form of longer access latency (additional delay cycles for the first request for data) relative to other types of RAM. Further, the problem of longer access latency is felt more sharply with shared DRAM, which needs to support concurrent operations required by the system, such as reading in new data from a DMU (data management unit) at the same time that a search for data in the memory is being performed.

[0005] In order to facilitate quicker storage and retrieval of data from the DRAM, a tree structure often is employed for the data being stored. For example, a typical tree structure may be from 12 levels to more than 23 levels deep. Such a large number of levels requires multiple requests to memory to obtain all of the necessary data, i.e., to access and utilize the desired leaf of the tree. In addition, with each successive level of the tree, there is more data (unsearched) than the previous level. These factors create further issues regarding how quickly traversal of a tree structure can occur.

[0006] Accordingly, what is needed is a system and method for optimization of a control structure for a leaf found from a tree search of data stored in external DRAM of an embedded processing system. The present invention addresses such a need.

Brief Summary of the Invention

[0007] Aspects for optimizing leaf comparisons from a tree search of data stored in external memory of an embedded processing system are described. The aspects include providing a control structure for leaf data comparisons as a control vector and a match key, and utilizing the control vector to direct types of comparison tests performed with the match key.

[0008] With the present invention, a leaf data control structure is provided that achieves a straightforward and efficient approach for improving leaf comparison operations of a tree search engine. These and other advantages of the present invention will be more fully understood in conjunction with the following detailed description and accompanying drawings.

Brief Description of the Several Views of the Drawings

- [0009] Figure 1 illustrates an overall block diagram of an embedded processing system.
- [0010] Figure 2 illustrates a search tree structure of PSCBs in accordance with the present invention.
- [0011] Figure 3 illustrates an example of a leaf data control structure in accordance with the present invention for a SMT algorithm.
- [0012] Figure 4 illustrates a layout of the example leaf data control structure of Figure 3 in a memory block.
- [0013] Figure 5 illustrates a layout of a leaf data control structure for FM/LPM algorithms in a memory block.
- [0014] Figure 6 illustrates a block diagram of key compare engines in accordance with the present invention.

Detailed Description of the Invention

- [0015] The present invention relates to control structures for tree searches in embedded processing systems. The following description is presented to enable one of ordinary skill in the art to make and use the invention and is provided in the context of a patent application and its requirements. Various modifications to the preferred embodiment and the generic principles and features described herein will be readily apparent to those skilled in the art. Thus, the present invention is not intended to be limited to the embodiment shown but is to be accorded the widest scope consistent with the principles and features described herein.
- [0016] The present invention presents aspects of providing optimal performance in a processing system utilizing shared RAM memories for both data and control storage. An overall block diagram of an embedded processing system applicable for utilization of the present invention is illustrated in Figure 1. As shown, the system 10 includes a central processing unit (CPU) core 12, the CPU core including a CPU 14, a memory management unit (MMU) 16, an instruction cache (I-cache) 18, and data cache (D-cache) 20, as is well appreciated by those skilled in the art. A processor local bus 22 couples the CPU core 12 to on-chip SRAM 24. Further coupled to the bus 22 is SDRAM (synchronous DRAM) controller 26, which is coupled to off-chip/external SDRAM 28.

A PCI (peripheral component interconnect) bridge 30 is also coupled to bus 22, the PCI bridge 30 further coupled to a host bus 32 that is coupled to host memory 34. As shown, a tree search engine 36 is also included and coupled to bus 22. The tree search engine 36 is a hardware assist that performs pattern analysis through tree searches to find the address of a leaf page for read and write accesses in the SDRAM 28.

[0017] In accordance with the present invention, the searches performed by the tree search engine 36 are improved with the optimization of a control structure for a leaf found from a tree search of data stored in external DRAM 28 of an embedded processing system. In general, tree searches, retrievals, inserts, and deletes are performed according to a key. Information is stored in the tree in leaves, which contain the keys as a reference pattern. To locate a leaf, a search algorithm processes input parameters that include the key pattern, and then accesses a direct table (DT) to initiate the walking of the tree structure through pattern search control blocks (PSCBs). The searches occur based on a full match (FM) algorithm, a longest prefix match (LPM) algorithm, or a software management tree (SMT) algorithm.

[0018] Figure 2 illustrates a search tree structure of PSCBs in accordance with the present invention and described in co-pending U.S. Patent Application, (docket no. RPS920020018US1/2493P) filed on November 22, 2002, serial no.10/065,819 and incorporated herein by reference in its entirety. By way of example, a search of the tree in Figure 2 begins with the memory access request of the left or right half of the Root or level 0 Branch Table (BT) based on the Next Bit Test (NBT) result from the Lookup Definition (LUDef) or Direct Table (DT, not shown) entry for this search tree. The access of the first branch table half contains multiple levels of PSCBs of the tree optimized for the search type. If after descending through the first table an external (lower) branch table address is arrived at instead of a leaf address, then an additional memory access request would be made for only the left or right half of this lower branch table. This process continues until a leaf address is arrived at during the decent through the lower branch table halves. When the search arrives at a leaf address, the process terminates with a memory access request for the leaf data to determine if a match was found.

[0019] Once the tree structure traversal leaf is located by a search, most hardware implementations will compare all of the key bits against the match key found in the leaf, and if a FM was required will report a failure if the key sizes or any key bits are different. If a LPM was required, then if the compare failed, the bit number from left to right of the first difference is returned to be used to select a matching prefix leaf (shorter key) passed during the decent phase of the tree search. If a SMT search was required, then the compare for equal must be modified to ignore the bits that are off in the mask fields (don't cares) and to change the compare to a magnitude range (min to max inclusive) in a number of other user defined fields in the key.

[0020] In accordance with the present invention, the comparison of the leaf data has better performance and flexibility with all information for key compare contained in each leaf (i.e., from a single memory reference) to reduce latency for match result. With the present invention, a control structure for leaf compares is provided that includes two major parts, a Control Vector and a Match Key. In order to demonstrate more fully the benefits and features of the control structure of the present invention, reference is made to an example shown in Figure 3 for an SMT leaf, since for comparison operations, it is the SMT algorithm that is primarily used for complex rules that may contain multiple don't care and/or magnitude range fields and is used to describe policy rules for security and quality-of-service types of applications. In a preferred embodiment, the control vector 2000 contains a 2-bit control setting for each byte of the key to be tested, e.g., 48 bits total for the 192-bit key (24 byte) implementation shown. These control settings are used to control the type of compare tests to be performed on each byte of the match key 2002.

[0021] The control settings along with the appropriate high (max) and low (min) bytes of the match key 2002 are processed from left to right so that magnitude comparisons of longer (L) ranges can be enabled by the compare results from the left or higher order bytes of the multi-byte range. The left most byte of a range compare is indicated by the Range (R) control setting and is then not dependent on the key byte to the left to enable the magnitude comparison. For simple masked (don't care) comparisons for equality the Mask (M) control setting is used and then the high and low match key bytes are used for the msk and val bytes respectively. The end of the match key is indicated by filling the rest of the control vector settings with the Exit or

Stop (x) value. In addition to terminating the comparison process for either a stop control or the maximum key length, the comparison may be terminated at the first failing byte of the key being tested.

[0022] In the example shown in Figure 3, the leftmost seven bytes, labeled 2004, of the 144-bit rule 2000 are defined as a masked compare (M control setting), the leftmost byte corresponding to byte 2006 of the match key. This is followed by an 8-bit range field, 2008, defined by the R control setting, which is then followed by two 16-bit range fields, 2010a and 2012a, each defined by the pairing of the R and L control setting, 2010b and 2012b. These ranges are then followed by an 8-bit mask field 2014 and then a 40-bit range 2016 which is indicated by the R followed by four Ls in the control vector 2000. This last range is then followed by Xs to fill out the rest of the 192-bit control vector which indicates the end of this rule.

[0023] The layout of this same example is shown in the SMT Leaf diagram of Figure 4. The control vector 2000 is contained in the six bytes shown as containing MMMM, MMMR, RLRL, MRLL, LLXX, and XXXX with their corresponding match key values indicated in correspondence with Figure 3. As shown, within the leaf following the control vector data and match key data, there is area available for additional data, like hash and encryption keys, sequence numbers, headers, protocols, etc., and as shown in the diagram, the portion of the match key space that is not required for a leaf's rule definition (XXXs) may be allocated as space for other additional data.

[0024] Application of the control structure and match key format to the FM and LPM leaf control blocks is shown in the FM/LPM leaf diagram of Figure 5. The Last Bit Tested (LBT) control byte 2020 indicates the right most bit of the match key (VALs) to be used for comparison, assuming the key is numbered from left to right and starts with zero. As in the SMT leaf block of Figure 3, there is area available for additional data, and the portion of the match key space that is not required for a leaf's rule definition (XXXs) may be allocated as space for other additional data.

[0025] With the SMT leaf's left to right and per byte approach to the control settings, a compare engine may be implemented in which any number of bytes of the key compare may be processed during a clock cycle. Figure 6 contains the block diagrams of key compare engines for both single and four byte examples. It should be obvious

put a limit on the maximum width of the engine.

[0027] Based on the cost versus performance approach for the SMT key compare engine, it should be obvious that a similar, incremental implementation would be appropriate for the FM/LPM compare engine, although a different width exclusive-or and priority encoder may be required to meet higher performance requirements, as is well appreciated by those skilled in the art.

[0028] In accordance with the leaf structure of the present invention, every SMT leaf can contain different mask and range field definitions. Further, SMT ranges from 1 to all bytes of key, and no separate range table is required in hardware for performance. Thus, the leaf structure of the present invention provides a straightforward and efficient approach for improving leaf comparison operations of a tree search engine.

[0029] Although the present invention has been described in accordance with the embodiments shown, one of ordinary skill in the art will readily recognize that there could be variations to the embodiments and those variations would be within the spirit and scope of the present invention. Accordingly, many modifications may be made by one of ordinary skill in the art without departing from the spirit and scope of the appended claims.